

Patent

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND APPARATUS FOR PREDICTION FOR
FORK AND JOIN INSTRUCTIONS IN SPECULATIVE EXECUTION**

INVENTOR:

**KEVIN W. RUDD
TIN-FOOK NGAI**

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD, SEVENTH FLOOR
LOS ANGELES, CA 90025-1030
(408) 720-8598

ATTORNEY'S DOCKET NO. 42P17887

Express Mail Certificate

"Express Mail" mailing label number: EV305339403US

Date of Deposit: December 8, 2003

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents, Alexandria, VA 22313-1450

Anne Collette

(Typed or printed name of person mailing paper or fee)

Anne Collette

(Signature of person mailing paper or fee)

12/8/2003

(Date signed)

METHOD AND APPARATUS FOR PREDICTION FOR FORK AND JOIN INSTRUCTIONS IN SPECULATIVE EXECUTION

FIELD

- 5 **[0001]** The present disclosure relates generally to microprocessors, and more specifically to microprocessors capable of speculative execution.

BACKGROUND

- 10 **[0002]** Modern microprocessors may support multi-threaded operation in their architectures. In some cases the multi-threaded operation may be sequential multi-threading, and in other cases the multi-threaded operation may be parallel multi-threading. In either case there are situations where a new thread may need to be spawned or where an existing thread may need to be merged back into the thread
- 15 that spawned it originally. The process of spawning a new thread may be called a fork operation, and the process of merging a thread back may be called a join operation. Fork and join operations may be coded in an operating system, or alternatively may be placed in executable code by the use of hardcoded fork and join instructions. The rationale
- 20 for using fork and join operations is to increase performance by the use of the forked-off threads. In some cases the forked thread may be part of non-speculative execution, but in other cases the forked thread may be speculative.

- 25 **[0003]** The use of hardcoded fork and join instructions may impact performance in several ways. If the instruction execution in the forked-off thread is correct and if the processor resources are not inadvertently impacted by the forked-off thread, then the performance may be

improved. However, if the instruction execution in the forked-off thread is incorrect, or if the processor resources are adversely impacted by the forked-off thread, then the performance may be reduced. It may be possible to consider the execution of a forked-off thread "desirable" in several different ways. It could be if the forked-off thread executed successfully. It could be if the overall processor execution throughput was enhanced. It could be a combination of these two, or it could take into account other measures of desirability.

[0004] Software execution could be used to determine whether it would be advantageous to take the fork or not. However, this determination would need to be accomplished prior to the fork, essentially occupying the resources available for both the main thread and the forked-off thread. The use of software determination of whether it would be advantageous to take the fork may use sufficient resources to impact processor performance by itself.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

5

[0006] **Figure 1** is a diagram showing the operation of a fork predictor supporting a speculative thread executing on a processor, according to one embodiment.

[0007] **Figure 2** is a schematic diagram of portions of a pipeline of a processor including a fork predictor, according to one embodiment.

10

[0008] **Figure 3** is a table showing operations in supporting and non-supporting processors, according to one embodiment of the present disclosure.

15

[0009] **Figures 4A and 4B** are a code fragment example and a flowchart of a method for speculative threads executing in a processor, according to one embodiment of the present disclosure.

[0010] **Figures 5A and 5B** are schematic diagrams of systems including a processor supporting execution of speculative threads, according to two embodiments of the present disclosure.

DETAILED DESCRIPTION

[0011] The following description describes techniques for a processor using multi-threaded execution to conditionally execute fork and join instructions without the use of extensive software testing prior to the execution of the fork. In the following description, numerous specific details such as logic implementations, software module allocation, bus signaling techniques, and details of operation are set forth in order to provide a more thorough understanding of the present invention. It will be appreciated, however, by one skilled in the art that the invention may be practiced without such specific details. In other instances, control structures, gate level circuits and full software instruction sequences have not been shown in detail in order not to obscure the invention. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate functionality without undue experimentation. In certain embodiments the invention is disclosed in the form of an Itanium ® Processor Family (IPF) processor or in a Pentium ® family processor such as those produced by Intel ® Corporation. However, the invention may be practiced in other kinds of processors that may wish to use conditional fork and join instructions in a multi-threaded environment.

[0012] Referring now to Figure 1, a diagram showing the operation of a fork predictor supporting a speculative thread executing on a processor is shown, according to one embodiment. In the Figure 1 embodiment, actions involving software occur on the right hand side of the figure while actions involving hardware occur on the left hand side of the figure. In other embodiments, the allocation of functions between software and hardware may be performed differently.

[0013] As execution proceeds in the main thread 110, a speculative fork instruction 122 may be reached. In various embodiments, speculative fork instruction 122 may be placed in the software by a compiler or by hand under a programmer's direction. Speculative fork instruction 122 may have the effect, if executed, of initiating the spawning of a speculative thread 120. The speculative fork instruction 122 may be under the control of a fork predictor 150, which may issue a prediction concerning any particular iteration of speculative fork instruction 122. When the speculative fork instruction 122 reaches the processor's execution units, the speculative fork instruction 122 may or may not be executed depending upon the prediction issued by fork predictor 150. If the prediction is that the speculative thread will be desirable, then speculative fork instruction 122 is executed, the main thread proceeds along 112 and in addition a speculative thread 120 is spawned. If, however, the prediction is that the speculative thread will not be desirable, then speculative fork instruction 122 is not executed, and the main thread proceeds along 112.

[0014] If the speculative thread 120 is spawned, at a certain later time it may merge back into the main thread 112. In one embodiment a join instruction 124 may be used to effect the joining of speculative thread 120 back into the main thread 112. In one embodiment, join instruction 124 may wait until both main thread 112 and speculative thread 120 have finished current processing before effecting the join. (Not shown are time-out exceptions for those cases where either the main thread 112 or speculative thread 120 are unable to finish current processing due to coding or system errors.) If the fork was executed

and successful, then the main thread 114, 116 may consume the results computed by the speculative thread 120.

[0015] Before the time that the main thread reaches the update instruction 126, it may be determined whether the execution of the speculative thread was successful. In one embodiment, “successful” may be equivalent to having executed correctly. This determination may take place in the main thread 110, 112, 114, in the speculative thread 120, or as part of a join instruction 124. If the execution of the speculative thread 120 was successful, then the main thread may progress without intervention. However, if the execution of the speculative thread 120 was not successful, then the results of the execution of the speculative thread 120 may be discarded and a recovery process may be initiated.

[0016] It may be possible to consider the forking of speculative thread 120 “desirable” in several different ways. It could be considered desirable if speculative thread 120 executed successfully, or would have executed successfully if it had been forked. It could be considered desirable if the overall processor execution throughput was enhanced, even in those cases where the associated execution of the speculative thread was determined not to be successful (e.g. the speculative thread could have advantageously made certain cache loads). It could be considered desirable if a combination of these two were present, or it could take into account other measures of desirability. This determination may be performed by an update instruction 126 or instructions, which in one embodiment may be a separate instruction and in another embodiment may be part of the join instruction 124. In either case, the update instruction 126 may send the results of the

determination of whether the execution of the speculative thread was desirable or not over an update signal path 164 to an update logic 152 in fork predictor 150. The update logic 152 may send these results on to the prediction logic 154 as part of the history information required by
5 prediction logic 154 to make predictions.

[0017] If the speculative thread 120 is not spawned, the main thread 112 may progress on by itself. In one embodiment, when the main thread 112 reaches the join instruction 124 it may treat the join instruction 124 as a no-operation (nop) in the absence of an executing
10 speculative thread. In this case, prior to the update instruction 126, the main thread 110, 112, 114 may determine whether the execution of the speculative thread 120 would have been desirable if it had been executed. This determination of being desirable may again be performed by an update instruction 126 or instructions, which in one
15 embodiment may be a separate instruction and in another embodiment may be part of the join instruction 124. In other embodiments, the determination of whether the execution of the speculative thread 120 would have been desirable may be made individually or separately and one determination may be made without the other. The update
20 instruction 126 may send the results of the determination of whether the execution of the speculative thread would have been desirable or not over an update signal path 164 to update logic 152 in fork predictor 150. The update logic 152 may send these results on to the prediction logic 154 as part of the history information required by prediction logic
25 154 to make predictions.

[0018] The prediction logic 154 may be one of various well-known branch predictor circuits adapted for use in predicting the outcome of

speculative threads. In differing embodiments, prediction logic 154 may implement prediction algorithms and methods well-known in the art such as a local history or a global history prediction algorithms. In other embodiments, other prediction algorithms specifically designed for speculative fork prediction may be used. One modification to such branch predictor circuits arises because the true results of a branch are automatically available to feed back into the predictor when the corresponding branch instruction is actually executed. Since the speculative thread 120 is a series of instructions, such feedback is not automatically available as the product of either the fork (where the prediction is made) or the join. This requires that the determination of whether the execution of the speculative thread was successful or not be separately made. In one embodiment, as mentioned previously, the determination may be performed by an update instruction 126 or instructions, which in one embodiment may be a separate instruction and in another embodiment may be part of the join instruction 124. The prediction logic 154 may be informed of the results of the determination by the update logic 152. In some cases, the update may never occur, and the prediction logic 154 may or may not take this fork prediction into account for subsequent predictions.

[0019] Referring now to Figure 2, a schematic diagram of portions of a pipeline of a processor 200 including a fork predictor 230 is shown, according to one embodiment. Instructions may be fetched or prefetched from a level one (L1) cache 202 by a prefetch/fetch stage 204. These instructions may be temporarily kept in one or more instruction buffers 206 before being sent on down the pipeline by an instruction dispersal stage 208. A decode stage 210 may take one or

more instructions from a program and produce one or more machine instructions. After exiting the decode stage 210, the instructions may enter the register rename stage 212, where instructions may have their logical registers mapped over to actual physical registers prior to execution. Upon leaving the register renaming stage 212, the machine instructions may enter a sequencer 214. In one embodiment, the sequencer 214 may schedule the various machine instructions for out-of-order execution based upon the availability of data in various source registers, or of data from some other source such as, for example, a results bypass network. In another embodiment, sequencer 214 may simply schedule the various machine instructions for in-order execution. Upon leaving the sequencer 214, the physical source registers may be read in register read file stage 216 prior to the machine instructions entering one or more execution units 218.

15 **[0020]** The execution units 218 may be configured to receive an input signal from a prediction logic 232 of a fork predictor 230. The execution of a speculative fork instruction may be conditioned by the prediction given by prediction logic 232: the speculative fork instruction may be executed if the prediction logic 232 predicts that the speculative thread

20 initiated by the speculative fork instruction will be desirable, and the speculative fork instruction may not be executed if the prediction logic 232 predicts that the speculative thread initiated by the speculative fork instruction will not be desirable. This behavior may be contrasted with a conditional non-speculative fork instruction, where the conditional

25 non-speculative fork instruction may or may not have its execution retired depending upon a predicate value which was externally determined and written but not affected by the fork prediction.

[0021] After execution in execution units 218, the results of the machine instructions may, in a retirement stage 220, update the machine state and write to the physical destination registers depending upon the resolved state of the corresponding predicate values. In one embodiment, a main thread may determine whether the execution of a speculative thread was desirable. This determination of desirability may be performed by an update instruction or instructions. The update instruction at retirement stage 220 may send the results of the determination of whether the execution of the speculative thread was desirable or not to an update logic 234 in fork predictor 230. The update logic 234 may send these results on to the prediction logic 232 as part of the history information required by prediction logic 232 to make predictions.

[0022] If the speculative thread had not been spawned, the main thread may progress on by itself. In one embodiment, prior to the time of an update instruction, the main thread may determine whether the execution of the speculative thread would have been desirable if it had in fact been executed. This determination may again be performed by an update instruction or instructions, which in one embodiment may be a separate instruction and in another embodiment may be part of the join instruction. The update instruction at retirement stage 220 may send the results of the determination of whether the execution of the speculative thread would have been desirable or not to update logic 234 in fork predictor 230. The update logic 234 may send these results on to the prediction logic 232 as part of the history information required by prediction logic 232 to make predictions.

[0023] The pipeline stages shown in Figure 2 are for the purpose of discussion only, and may vary in both function and sequence in various processor pipeline embodiments. Similarly the connections of fork predictor 230 to the pipeline stages may vary in various embodiments.

- 5 In some embodiments fork predictor 230 may be part of a branch prediction circuit. In other embodiments, fork predictor 230 may be implemented as an independent structure.

- [0024]** Referring now to Figure 3, a table showing operations in supporting and non-supporting processors is shown, according to one
10 embodiment of the present disclosure. In order to facilitate use of common software in processors of differing configurations, the common software should be capable of executing on supporting processors that implement a speculative fork instruction and also on non-supporting processors that do not. In one embodiment a non-supporting processor
15 could trap on the non-implemented speculative fork instruction and handle the situation as an exception. However, this would consume valuable time and processor resources. Therefore, in another embodiment the software should test for the presence or absence of support prior to attempting the execution of the speculative fork
20 instruction.

- [0025]** Figure 3 shows one embodiment of three possible cases when executing a code sequence (such as that shown in Figure 4A below) in either a supporting or a non-supporting processor. In the non-supporting processor, the speculative fork, join, and update
25 instructions may be executed as nops. In the supporting processor, whether or not the speculative thread was in fact forked, the determination of whether the speculative thread was or would have

been desirable may be made and the update to the fork predictor may be made.

[0026] Referring now to Figure 4A, a code fragment example is shown, according to one embodiment of the present disclosure. The code
5 fragment example has a main routine that performs two sets of computations: the computation that will be performed in a main thread in parallel with a speculative thread and the post-speculation computation that consumes the results of the earlier computation. The speculative computation may simplify the main computation and may
10 enable execution of the post-speculation computation much earlier than would be possible with the non-speculative form of this code, thus shortening the critical path through this code in the event that the speculation is successful.

[0027] Referring now to Figure 4B, a flowchart of a method for
15 speculative threads executing in a processor is shown, according to one embodiment of the present disclosure. The process 400 may begin with a test instruction 410. The test instruction 410 may in various embodiments consist of reading a processor identification register or may consist of a specialized test instruction. After executing the test
20 instruction 410 the process enters decision block 412, where it may be determined whether or not the processor supports the speculative fork instruction. If not, the process exits via the NO path and exits the process. If so, then the process exits via the YES path and enters decision block 420.

[0028] In decision block 420 it may be determined whether the fork predictor issues a prediction that it would be desirable to execute a speculative thread. If not, then the process exits via the NO path, and

in block 434 determines whether the speculative thread would in fact have been desirable before updating the fork predictor with this result in block 446. If, however, the fork predictor issues a prediction that it would be desirable to execute the speculative thread, the process exits via the YES path and in block 424 spawns off a speculative thread via a speculative fork instruction before entering decision block 428.

[0029] In decision block 428 it may be determined if both the main thread and the speculative thread have executed to completion before a join instruction. If not, then the process exits via the NO path and decision block 428 repeats. When both the main and speculative threads have executed to completion, then the process exits via the YES path and in block 430 the join instruction is executed before entering decision block 438. (Not shown are time-out exceptions for those cases where either the main thread or speculative thread is unable to execute to completion due to coding or system errors.)

[0030] In decision block 438 it may be determined whether the speculative thread was executed successfully. If so, then the process exits via the YES path, and in block 434 determines whether the speculative thread would in fact have been desirable before updating the fork predictor with this result in block 446. If not, then the process exits via the NO path and initiates a recovery in block 442. Then in block 434 the process determines whether the speculative thread would in fact have been desirable before updating the fork predictor with this result in block 446.

[0031] An example of instructions that may be used in the Figure 4B process executing on an Itanium ® family or compatible processor is as follows. For the test block 410, it may be possible to use a “test feature”

tf instruction, to gate the speculative fork instruction and thereby avoid trapping and exception handling on non-supporting processors. The code line

```
tf    p1, p0 = @fork
```

- 5 would set the predicate p1 true on supporting processors and false on non-supporting processors. If speculative forking is available, then the modified branch instruction br.fork could be used as the speculative fork instruction as in the code lines

```
forkpoint:
```

- 10 (p1) br.fork.spec speculative_routine

whose execution would be suppressed if predicate p1 is false. (The speculative fork may be made conditional by merging another condition with the br.fork.spec qualifying predicate p1. For example, (p1)

cmp.unc p1, p0 = r0, r2. This may ensure that if p1 is already false

- 15 that it will propagate to its result and force it to be false as well.) The join instruction could be implemented using a nop with hint which is just a nop if the @join hint is not supported. This could be expressed as in the code lines

```
join_point:
```

- 20 nop.hint @join

for use in the Figure 4B process. In other embodiments, the join instruction could be implemented as a new instruction in branch predict/hint space where undefined operations are ignored. The determination of whether the speculative thread was or would have
25 been desirable could be implemented with a compare/equal/unconditional cmp.eq.unc instruction as in the following code line:

(p1) `cmp.eq.unc p0, p3 = r0, r1`

where the `cmp` instruction would check the result of the desirability computation in register `r1` and sets predicate `p3` true if the speculative thread was or would have been desirable and false otherwise. A similar
5 compare could determine whether the speculative thread was successfully executed. In other embodiments, other code sequences that compute whether or not the speculative thread was or would have been successful could be used. The update instruction could be implemented as a modified branch predict `brp.join` instruction as in the
10 following code line:

(p3) `brp.join speculative_routine, fork_point`

which is ignored if not implemented because it is in a portion of branch prediction/hint space where undefined operations are simply ignored.

[0032] Referring now to Figures 5A and 5B, schematic diagrams of
15 systems including a processor supporting execution of speculative threads are shown, according to two embodiments of the present disclosure. The Figure 5A system generally shows a system where processors, memory, and input/output devices are interconnected by a system bus, whereas the Figure 5B system generally shows a system
20 where processors, memory, and input/output devices are interconnected by a number of point-to-point interfaces.

[0033] The Figure 5A system may include several processors, of which only two, processors 40, 60 are shown for clarity. Processors 40, 60 may include level one caches 42, 62. The Figure 5A system may
25 have several functions connected via bus interfaces 44, 64, 12, 8 with a system bus 6. In one embodiment, system bus 6 may be the front side bus (FSB) utilized with Pentium® class microprocessors manufactured

by Intel® Corporation. In other embodiments, other buses may be used. In some embodiments memory controller 34 and bus bridge 32 may collectively be referred to as a chipset. In some embodiments, functions of a chipset may be divided among physical chips differently
5 than as shown in the Figure 5A embodiment.

[0034] Memory controller 34 may permit processors 40, 60 to read and write from system memory 10 and from a basic input/output system (BIOS) erasable programmable read-only memory (EPROM) 36. In some embodiments BIOS EPROM 36 may utilize flash memory or
10 other memory devices. Memory controller 34 may include a bus interface 8 to permit memory read and write data to be carried to and from bus agents on system bus 6. Memory controller 34 may also connect with a high-performance graphics circuit 38 across a high-performance graphics interface 39. In certain embodiments the high-
15 performance graphics interface 39 may be an advanced graphics port AGP interface. Memory controller 34 may direct read data from system memory 10 to the high-performance graphics circuit 38 across high-performance graphics interface 39.

[0035] The Figure 5B system may also include several processors, of
20 which only two, processors 70, 80 are shown for clarity. Processors 70, 80 may each include a local memory channel hub (MCH) 72, 82 to connect with memory 2, 4. Processors 70, 80 may exchange data via a point-to-point interface 50 using point-to-point interface circuits 78, 88. Processors 70, 80 may each exchange data with a chipset 90 via
25 individual point-to-point interfaces 52, 54 using point to point interface circuits 76, 94, 86, 98. Chipset 90 may also exchange data with a high-

performance graphics circuit 38 via a high-performance graphics interface 92.

[0036] In the Figure 5A system, bus bridge 32 may permit data exchanges between system bus 6 and bus 16, which may in some
5 embodiments be a industry standard architecture (ISA) bus or a peripheral component interconnect (PCI) bus. In the Figure 5B system, chipset 90 may exchange data with a bus 16 via a bus interface 96. In either system, there may be various input/output I/O devices 14 on the bus 16, including in some embodiments low performance graphics
10 controllers, video controllers, and networking controllers. Another bus bridge 18 may in some embodiments be used to permit data exchanges between bus 16 and bus 20. Bus 20 may in some embodiments be a small computer system interface (SCSI) bus, an integrated drive electronics (IDE) bus, or a universal serial bus (USB) bus. Additional
15 I/O devices may be connected with bus 20. These may include keyboard and cursor control devices 22, including mice, audio I/O 24, communications devices 26, including modems and network interfaces, and data storage devices 28. Software code 30 may be stored on data storage device 28. In some embodiments, data storage device 28 may
20 be a fixed magnetic disk, a floppy disk drive, an optical disk drive, a magneto-optical disk drive, a magnetic tape, or non-volatile memory including flash memory.

[0037] In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will,
25 however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and

Assignee: Intel Corporation

drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.